Journal of Experimental Psychology: General
1997, Vol. 126, No. 1, 54–70

0096-3445/97/$3.00

Bidding on the Future: Evidence Against Normative Discounting of Delayed Rewards

Kris N. Kirby
Williams College

Normative economic models assume that delay-discounting rates for future rewards are independent of the amount of, and delay to, a reward. These assumptions were tested in 3 experiments in which participants were asked to bid their own money on delayed monetary rewards in a sealed, second-bid auction. In Experiments 1 and 2, participants made their bids without feedback about the bids of other participants. In Experiment 3, half of the participants received such feedback. In Experiments 2 and 3, participants adjusted their bids until they were indifferent between the bid and the delayed reward. Participants' bids violated both normative assumptions: Discounting rates decreased with increases in amount for 62 of 67 participants, and a function in which the rate decreases with increases in delay (hyperbolic) fit the bids better than did a normative function (exponential) for 59 of 67 participants.

A future reward is typically worth less to us than the same reward available immediately. This *delay discounting* of individual future rewards is intuitively obvious and easily summarized: the sooner, the better. However, not so obvious is the quantitative form of the function by which the present value of a reward is discounted as the delay to the reward increases. Normative economic models assume that present value decreases by a fixed proportion per unit of time that one must wait for the reward, that is, that future value is discounted *exponentially* with delay (e.g., Fishburn & Rubinstein, 1982; Lancaster, 1963; Meyer, 1976). Just as a bank account increases by compounding a fixed rate of interest over time, the present value of a future reward grows as that reward approaches by compounding a fixed rate of increase over time. In contrast, empirical work over the past 35 years has suggested that value may increase by an increasingly larger proportion per unit time as the reward approaches (e.g., Baum & Rachlin, 1969; Killeen, 1970; Logan, 1960; Mazur, 1987). This innocuous-sounding distinction turns out to be quite important because under the normative model a reward that is preferred to another from one temporal vantage point is preferred from any temporal vantage point. However, if the discounting function deviates from the normative model, it is possible for preferences to reverse over time. Such discounting-related preference reversals offer an elegant explanation of a wide range of suboptimally impulsive behavior, from the scratching of itches to substance abuse (Ainslie, 1975, 1992; Herrnstein, 1990; Rachlin, 1974; Rachlin & Green, 1972). Our preference to resist temptation weakens and then succumbs as the object of that temptation draws temporally closer.

Figure 1 illustrates normative discounting and two types of violations of normative discounting that can lead to impulsive preference reversals. The top panel shows normative discounting functions for two delayed rewards—a smaller, earlier reward (S) and a larger, later reward (L)—with their undiscounted values at the time of receipt indicated by the heights of the vertical lines. From left to right in the figure the increasing discounting curves show how the present values of the two rewards increase as they approach in time. In the top panel, the discounting curves are exponential in form, and the present values of the two rewards maintain a fixed ratio, with the smaller reward preferred from all temporal vantage points. In simpler terms, these curves never cross. Normative, exponential discounting can be expressed as:

$$V = Ae^{-kD}, \quad (1)$$

where V is the present, discounted value of the delayed reward, A is the amount of the delayed reward, D is the delay until receipt of the reward, and k is the discounting rate. (All delays described in this paper are measured in days, and the values of k are scaled accordingly.)

In Equation 1 the rate at which future rewards are discounted is independent of the delay to the reward (*rate–delay independence*). The second panel in Figure 1 shows the same larger, later reward as the one in the top panel (the black curve), and also shows how this curve would change if the functional form were *hyperbolic* (the gray curve), a type of function that violates rate–delay independence. Essentially, the hyperbolic function is relatively steeper at short delays and flatter at long delays, corresponding to the decrease in the discounting rate as the delay increases. This increased "bend" in the curve can allow the curves for two different rewards to cross, as illustrated by the two hyperbolic curves in the third panel in Figure 1. As time passes, this person's preference reverses from preference for the larger reward at long delays to preference for the smaller

I thank Jonathan Baron for thoughtful comments on the article, David Laibson for introducing me to second-bid auctions, and Todd Poret for conducting the auctions and assisting with data analyses.

Correspondence concerning this article should be addressed to Kris N. Kirby, Bronfman Science Center, Williams College, Williamstown, Massachusetts 01267. Electronic mail may be sent to kris.kirby@williams.edu.

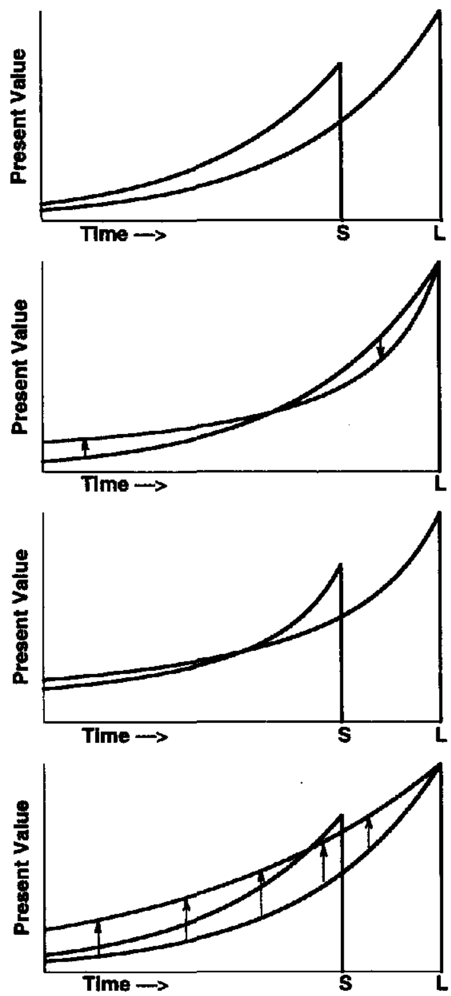

Figure 1. Present value of two delayed rewards as a function of delay. *S* indicates the point of receipt of a smaller, earlier reward, and *L* indicates the point of receipt of a larger, later reward. In the first panel, both rewards are discounted at the same exponential rates, yielding curves that do not cross. In the second panel, the arrows indicate how the original exponential curve for the larger reward (black) would change if the discounting rate were inversely related to delay (gray). In the third panel, the discounting rates for both rewards are inversely related to delay, yielding curves that cross. In the fourth and bottom panel, the arrows indicate how the original exponential curve for the larger reward (black) would "flatten" if its discounting rate were decreased (gray), yielding curves that cross.

reward at short delays. Suppose, for example, that one were asked to choose between spending a night on the town a week from next Monday or having a productive day at work a week from next Tuesday. At this delay one might prefer the productive day's work. However, as that Monday approaches such preferences may reverse, so that if one had the opportunity to change one's mind, say on Monday afternoon, one impulsively might choose the night on the town. The discounting model of impulsiveness hinges on the possibility that the discounting curves for different future rewards can cross. This can happen with hyperbolic functions, such as the following (Mazur, 1987):

$$V = \frac{A}{1 + kD}. \tag{2}$$

Hyperbolic discounting functions are implied by the empirically well-established matching law (Davison & McCarthy, 1988; De Villiers & Herrnstein, 1976), which predicts that the discounting rate is inversely related to delay (Ainslie & Herrnstein, 1981; Chung & Herrnstein, 1967; Herrnstein, 1981; Mazur & Herrnstein, 1988).[1]

In addition to rate–delay independence, normative models also assume that the discounting rate (k in Equation 1) is independent of the size of the reward (*rate–amount independence*). That is, a large reward must be discounted by proportionally the same amount per unit time as a smaller reward. This is true, for example, of the two rewards in the top panel of Figure 1. Deviations from this assumption could also lead to crossing discounting curves, as pointed out by Green, Fisher, Perlow, and Sherman (1981). Specifically, if the discounting rate were inversely related to the size of the reward, then discounting curves could cross even if each individual curve were exponential in form. This is illustrated in the bottom panel of Figure 1. The curve for the smaller reward and the lower curve for the larger reward (in black) are identical to those shown in the top panel. The arrows indicate how the curve for the larger reward would change (shown in gray) if the discounting rate for that reward were decreased. The curve flattens out, allowing it to cross the curve for the smaller reward, even though both curves meet the rate–delay independence assumption (i.e., are exponential).

Exponential discounting with a fixed rate across reward sizes is normative, in the prescriptive sense, precisely because it would preclude preference reversals solely due to changes in temporal vantage point.[2] Behavior that is influenced by such preference reversals may exhibit systematic, "dynamic inconsistencies" over time (Fishburn & Rubinstein, 1982; Koopmans, 1960; Loewenstein, 1992; Machina, 1989; Prelec, 1989; Strotz, 1955), leading to both wasteful planning and impulsive, suboptimal choices. For example, a planner who makes decisions based on preferences at early stages of planning may later reverse preferences and make

[1] I refer to k in Equation 2 as the *hyperbolic discounting rate parameter* to avoid confusion with the discounting rate per se, by which I mean the percentage change in value per unit time as indexed by k in Equation 1. It is the latter, the percentage change per unit time, that is assumed to be constant in rate–delay independence. The hyperbolic rate parameter is constant across delay even when the percentage change is inversely related to delay.

[2] Note that there is nothing irrational about discounting per se. For example, it may be economically sensible for a person to accept a smaller monetary amount now and put it in a bank where it can earn interest rather than wait to obtain the same or larger amount at a future date. It also can be adaptive for an animal to choose an immediate prey if it risks starvation before the next opportunity to obtain that same or larger prey (Houston & McNamara, 1985; Logue, 1988).

decisions that undo earlier plans (Pollak, 1968; Prelec, 1989). The myopia of such inconsistency can be illustrated with an example from Herrnstein (1990). People sometimes report that (a) they would prefer to receive $115 in 53 weeks over $100 in 52 weeks but also report that (b) they prefer $100 today over $115 in 1 week. Choosing to wait for the $115 reward in Scenario (a) is intuitively reasonable. The trouble is that after a year passes the individuals in Scenario (a) will be in precisely the same relation to the two rewards as they are in Scenario (b), in which they said they prefer not to wait for the larger reward. Although some notion of "rationality" might be preserved in this inconsistency, such unstable preferences are not easily reconciled with a descriptive model based on normative assumptions (Herrnstein, 1990).

Such preference reversals do occur. Evidence consistent with systematic preference reversals due to crossing discounting curves has accumulated from studies using pigeons (Ainslie & Herrnstein, 1981; Green et al., 1981; Navarick & Fantino, 1976; White & Pipe, 1987), rats (Boehme, Blakely, & Poling, 1986; Deluty, 1978; Logan & Spanier, 1970), adult humans (Ainslie & Haendel, 1983; Kirby & Herrnstein, 1995; Millar & Navarick, 1984; Navarick, 1982; Solnick, Kannenberg, Eckerman, & Waller, 1980; Winston & Woodbury, 1991), and children (Burns & Powers, 1975). For example, in the three experiments reported in Kirby and Herrnstein (1995), participants were offered choices between smaller, earlier rewards and larger, later rewards, and the delays to those rewards were adjusted on the basis of participants' choices to determine whether preferences would reverse systematically with delay. Real money and goods ranging in value from $12 to $52 were used as rewards. Thirty-four of the 36 participants consistently chose the smaller reward when the delays to both rewards were short, but reversed preference and consistently chose the larger reward as the delays to both rewards were lengthened. Additional data consistent with preference reversals come from experiments in which subjects were allowed to "precommit" themselves to the larger reward when the delays to both rewards were long. When given this opportunity to deny themselves the option of an impulsive choice, pigeons (Ainslie, 1974; Rachlin & Green, 1972), rats (Deluty, Whitehouse, Mellitz, & Hineline, 1983), and humans (Solnick et al., 1980) frequently did so. This behavior makes sense if the animal and human subjects anticipated preference reversals as the smaller rewards approached in time. On trials in which precommitment was not possible, the same animal and human subjects typically chose the smaller reward when the delay to that reward was short.

In sum, the data on preference reversals with delay are consistent with crossing discounting curves, as illustrated in either of the bottom two panels in Figure 1. However, none of these preference-reversal studies tell us whether violations of the rate–amount assumption, the rate–delay assumption, or both are responsible for the deviation from the normative model. Thus, it becomes important to determine whether discounting curves do deviate from normative models, and if so, precisely how they deviate.

Evidence Bearing on Rate–Delay Independence

A number of studies using animal subjects have provided evidence consistent with an inverse relationship between discounting rate and delay (e.g., Baum & Rachlin, 1969; Chung & Herrnstein, 1967; King & Logue, 1992; Logan, 1965; Logan & Spanier, 1970; Mazur, 1984, 1994). Responding is typically consistent with the generalized matching law (Baum, 1974), which, as noted above, implies hyperbolic discounting (but see, e.g., Killeen, 1985, for possible deviations from matching). However, those studies were not designed to directly compare hyperbolic and exponential functions. Mazur (1987) and Rodriguez and Logue (1988, Experiment 1) did design experiments to make such a comparison. Using discrete-choice procedures, the delays to a smaller, earlier reward (2-s access to food) and a larger, later reward (6-s access to food) were adjusted based on pigeons' choices until the pigeons were indifferent between the two rewards. At this indifference point the present values of the two rewards are equal, and the hyperbolic function in Equation 2 predicts that the delay to a larger reward of amount A_L will be a linear function of the delay to a smaller reward of amount A_S, with a slope equal to A_L/A_S, which must be greater than 1. The exponential function in Equation 1 also predicts a linear relationship, but with a slope equal to 1. Both studies found that the slope was greater than 1, consistent with the predictions of the hyperbolic function. Unfortunately, as pointed out by Green and Myerson (1993), these slope predictions tacitly assume rate–amount independence. If the discounting rate is inversely related to amount, then Equation 1 also predicts a slope greater than 1 (see Green, Fristoe, & Myerson, 1994, for a detailed discussion). Thus, Mazur's, and Rodriguez and Logue's, experiments provide evidence that at least one of the normative assumptions is violated, but they do not tell us which.

The data from studies with humans have a variety of limitations. At least six studies using hypothetical rewards have found that the discounting rate is inversely related to delay, in violation of the rate–delay independence assumption (Ainslie & Haendel, 1983; Benzion, Rapoport, & Yagil, 1989; Loewenstein, 1987; Rachlin, Raineri, & Cross, 1991; Thaler, 1981; Winston & Woodbury, 1991). However, the use of hypothetical rewards limits the generalizability of the results. In addition, each of these studies inferred the inverse relationship from data aggregated across participants, which can introduce an artifactual bias favoring hyperbolic functions. To illustrate how this can arise, suppose that two people both discounted delayed rewards exponentially in line with normative assumptions, but that Person A had an exponential discount rate of 5% per day and Person B had a rate of 50% per day (the artifact does not depend on using such extreme rates, but they aid the illustration). With a 1-day delay, a $100 reward would be worth $95 to Person A and $50 to Person B. With a 2-day delay, the $100 would be worth $90.25 to Person A and $25 to Person B. The artifact arises when one averages the present values of the two people. At the 1-day delay, the average is $72.50, which represents a 28% decrease. At the 2-day delay, the average is $57.63, which represents a 21% decrease from the 1-day mean. Thus, the mean data appear to show that the rate is inversely related to delay, but this trend does not reflect any individual person's underlying discounting curves. Because group averaging can both introduce artifactual biases and obscure genuine relationships (cf. Estes, 1956; Nosofsky, Palmeri, & McKinley, 1994), it is imperative that the inverse rate–delay relationship be demonstrated at the individual level.

A small number of studies with humans that used real rewards have found evidence for an inverse rate–delay relationship, but each of these also has important limitations. Rodriguez and Logue (1988, Experiment 2) adjusted delays to rewards in the same manner as in their and Mazur's (1987) experiments with pigeons and reported evidence for hyperbolic functions from the observed indifference points. As with the pigeon studies, however, the inferred support for hyperbolic discounting tacitly assumed rate–amount independence. In addition, Rodriguez and Logue used as the delayed rewards points that were exchangeable for money at the end of the experimental session. One difficulty with using points as rewards, as discussed by Logue, Peña-Correal, Rodriguez, & Kabela, (1986), is that there is little advantage to choosing the more immediate points because they cannot be exchanged for money until the end of the session anyway. In fact, experiments with humans using points as rewards rarely find choices of the smaller reward (Belke, Pierce, & Powell, 1989; Logue, King, Chavarro, & Volpe, 1990; Logue et al., 1986) except when such choices yield greater reward densities (Flora, 1995; Flora & Pavlik, 1992) or greater total reward at the end of the session (Logue et al., 1990). Rewards that can be consumed during the experimental session, such as access to video game playing (Millar & Navarick, 1984), termination of an aversive noise (Navarick, 1982; Solnick et al., 1980), and access to juice (Forzano & Logue, 1992, 1995; Logue & King, 1991) lead to substantially more within-session choices of the smaller of two rewards than do points exchangeable for money or points exchangeable for food (Forzano & Logue, 1994). Hyten et al. (1994) have shown that manipulating delays to acquiring points produces little change in preferences, but that manipulating the delays to when those points are exchangeable for money does affect preferences. Together, these difficulties with using delayed points exchangeable for money raise a question as to whether Rodriguez and Logue's (1988) evidence against exponential discounting in humans (even assuming rate–amount independence) would generalize to choices that did not involve the use of points exchangeable for money.

Holcomb and Nelson (1992) offered participants choices between smaller, earlier rewards and larger, later rewards in a questionnaire format. One trial was chosen at random from the questionnaire, and each participant received the chosen reward on that trial. A small reward was paired with multiple larger, later rewards, with the amounts for the larger rewards generated by increasing the amount of the smaller reward by a fixed rate per additional day of delay. For example, in one condition the larger reward was generated by increasing an earlier $5 reward by 1.5% per day of additional delay. So for a 1-day delay the larger, later reward was $5.07, and for the 14-day delay the larger reward was $6.15. If discounting were exponential, then a person who preferred $5 now over $5.07 tomorrow should also prefer $5 now over $6.15 in 14 days. Instead, Holcomb and Nelson found that the percentage of participants who chose the smaller, earlier reward decreased as the delay (and size) of the larger, later reward increased. They concluded that this was evidence for hyperbolic discounting. However, this conclusion tacitly assumes rate–amount independence. Because the later rewards were increased in amount as their delays were increased, it is possible that those rewards were being discounted at lower rates, which could account for the increase in choices of the larger rewards.

To determine the precise form of a discounting function, one ideally would like to estimate people's true present values of a real delayed reward at a number of delays. (I am calling the "true" value that amount for which the individual would be indifferent between receiving that amount immediately and receiving the delayed reward.) However, there is a problem with asking participants to report immediate amounts corresponding to their true values of delayed rewards when they may actually receive one of the immediate amounts that they report. It is in their interest to report amounts as high as possible, rather than report their true values, so as to earn as much money as possible in the experiment. To offset this disincentive for accuracy, Kirby and Maraković (1995) used a (simulated) *sealed, low-bid-wins auction.* Participants were asked to bid the least amount of money that they would be willing to accept immediately in exchange for receiving the delayed rewards, without knowledge of the bids of other participants. Because overbidding one's true value of the delayed reward reduces one's chances of winning the bid, the incentive for overbidding is reduced. Both hyperbolic and exponential functions were fit to the participants' bids separately for each delayed reward amount (thereby avoiding any confounding with a possible rate–amount relationship). For 20 of 21 participants in that experiment, and for every delayed reward amount that was offered, the hyperbolic function in Equation 2 provided a better account of the data than did the exponential function in Equation 1.

Unfortunately, although the low-bid-wins procedure reduces the incentive for overbidding, it does not eliminate it. Losing the bid means receiving the delayed reward, and winning the bid at one's true value means receiving a reward of equal value immediately. So long as participants bid at least as high as their true values they cannot receive a reward of less value to them than the delayed reward. But by overbidding, a participant opens the possibility of winning an immediate reward of greater value than the delayed reward, although with decreasing probability as the (over) bid increases. Because the comparisons between hyperbolic and exponential fits were relative, the better fit for the hyperbolic function is not compromised by the possible upward bias in discounting rates caused by the low-bid-wins procedure discussed earlier as long as that bias is constant across delays. However, it is plausible that such a bias might increase as bids get lower. If one assumes a linear decrease in the probability of winning as one's bid increases, then the difference between the bid with the maximum expected value and one's true value of the delayed reward increases as one's true value decreases. Concretely, overbidding one's true value by $1 would cause a greater decrease in the probability of being the low bid when one's true value is, for example, $18 than when it is $8. Because present value decreases as delay increases, the upward bias could also increase with delay, and this would artificially favor the hyperbolic function. Because of this possibility it is important to replicate the rate–delay relationship using a procedure that does not have this limitation.

One study that did not have this bias was Horowitz (1991, Conditions 1 and 2), in which students in a classroom setting were offered a chance to bid their own money on $50 bonds in a *sealed, first-rejected-price auction.* For example, in Condition 1 the top four bidders were allowed to purchase the bonds for the price bid by the fifth-highest bidder (the first rejected price). This type of auction eliminates any incentive for over- or underbidding. Separate auctions were conducted for bonds that matured in either 34 or 64 days. Not only were the observed discounting rates inversely related to delay, but the average bid (mean or median) for the 34-day delay was actually lower than the average bid for the 64-day delay. This implies a discounting curve with a U shape, which would be difficult to reconcile with any discounting model. However, this study had a number of important limitations that caution against taking this result too seriously. First, the difference in discounting rates was inferred from the difference in average bids between the 34- and 64-day delay conditions, rather than computed as a within-subject effect. A moderate rank-order correlation between participants' bids in the two conditions ($\rho = .45$) raises the possibility that the U shape could be an artifact of group averaging. Second, only two of the six winners of the auctions were willing to purchase the bonds that they had bid for, suggesting that the bids did not reflect their true values. Finally, the same group of participants was tested in both conditions, with the 34-day condition being conducted 1 month after the 64-day condition. Thus, order effects were not controlled, and during the interval between auctions some of the participants' discounting rates could have changed. Therefore, the results of this experiment are problematic, and in any case, they cannot be taken as evidence for any form of monotonically declining discounting function, exponential or hyperbolic.

Evidence Bearing on Rate–Amount Independence

Several studies have examined differences in the sensitivity of behavior to changes in amount and delay using real rewards with pigeons (Green et al., 1981; Green & Snyderman, 1980; Logue, Rodriguez, Peña-Correal, & Mauro, 1984; Rodriguez & Logue, 1986), rats (Ito, 1985; Ito & Asaki, 1982; Tobin, Chelonis, & Logue, 1993), and humans (Logue, Forzano, & Tobin, 1992). However, none of these studies examined the relationship between the discounting rate and the amount of the reward. Of the several published studies that have directly examined this relationship with human participants, most have found evidence for an inverse rate–amount relationship, consistent with the discounting curves in the bottom panel of Figure 1. However, the majority of these either used hypothetical rewards (Benzion et al., 1989; Green, Fristoe, et al., 1994; Green, Fry, & Myerson, 1994; Kirby & Maraković, 1995, Experiment 2; Raineri & Rachlin, 1993; Thaler, 1981) or chose participants by lottery to receive a real reward (Kirby & Maraković, 1996). Unfortunately, the only two studies in which each participant received a real reward found conflicting results. Holcomb and Nelson (1992), discussed above, compared participants' choices across pairs of smaller, earlier rewards and larger, delayed rewards for which the ratio of the two rewards was held constant but the absolute magnitudes of the two rewards were varied. For example, in one pair the choice was between $6.15 in 15 days and $5 tomorrow (a larger–smaller ratio of 1.23); in another pair the choice was between $20.94 in 15 days and $17 tomorrow (also a larger–smaller ratio of 1.23). If discounting rates are independent of amount, both Equations 1 and 2 predict that participants choose either the smaller rewards in both pairs or the larger rewards in both pairs. In fact, Holcomb and Nelson found that a higher percentage of participants chose a larger reward when the absolute magnitudes were increased. But because they did not estimate discounting rates for the different reward sizes for each participant, and reported only aggregate data, one cannot determine how many participants actually showed this effect. Nonetheless, the data are inconsistent with rate–amount independence.

The low-bid-wins auction conducted by Kirby and Maraković (1995, Experiment 1) also directly examined rate–amount independence, but did so within subjects so as to avoid any artifacts of group averaging. In that experiment, in which each participant received a real reward, there was no evidence that the discounting rate varied with the amount of the reward. (This lack of a relationship cannot be dismissed as due to a lack of power. Based on the effect sizes found in previous studies, this experiment had power greater than .90 at the $p = .05$ level.) One possible explanation for the failure to find an inverse rate–amount relationship is that a floor effect on discounting rates may have been encountered due to a possible biasing artifact of the procedure. As discussed above, the low-bid-wins auction does not completely remove the incentive for overbidding, and any resulting bias upward in bidding could have biased the estimated discounting rates downward, thereby making the discounting rates across the reward sizes more difficult to distinguish.

In summary, previous data on the inverse relationship between discounting rates and the length of the delay to the reward have tended to favor hyperbolic over exponential discounting. However, none of them strongly compel rejection of rate–delay independence. The evidence bearing on rate–amount independence is inconsistent, with data from hypothetical and probabilistic rewards indicating an inverse relationship between rate and amount, but data from the two studies with humans that have used real rewards yielding conflicting results. The purpose of the experiments reported below was to test both assumptions using a new procedure and an analytical method that together avoid the limitations of previous research.

General Method

To encourage participants to report accurate present valuations of delayed rewards, they participated in a series of *sealed, second-bid auctions* in which they placed bids indicating the highest amounts that they were willing to pay immediately to receive delayed monetary rewards in the number of days specified. In a second-bid auction the highest bid wins, but instead of purchasing the delayed reward at the price of his or her own bid, the winner is allowed to purchase the item at the price of the second-highest bid. (The second-bid auction is a special case of first-rejected-price auction.) Because purchasing the delayed reward at a price below one's true value represents a profit, in this type of auction the optimal strategy is to bid one's true value (Vickrey, 1961). There is no advantage to bidding below one's true value because bidding low cannot decrease the amount one pays (which is determined by another bidder), it can only reduce one's chance of winning at a profit. There is also no advantage to bidding above one's true value because this will only increase one's odds of winning when another bidder has also bid above one's true value, in which case winning would require one to pay more than the item is subjectively worth. For the goals of the present research this creates an *incentive compatible procedure;* participants need not concern themselves with the probability of winning, but instead may focus on submitting bids that accurately reflect their discounted values of the delayed rewards. These auctions were not simulated. Participants bid their own money for the option of purchasing real delayed rewards.

Bids were analyzed using nonlinear regression to fit Equations 1 and 2 to each participant's bids for each delayed reward amount separately. By fitting these functions within subjects, one avoids the possible artifacts that can arise in aggregate data and can obtain an estimate of each participant's discounting rate. By fitting the curves for each delayed-reward amount separately one can assess the relative fits of the hyperbolic and exponential functions without tacitly assuming rate–amount independence and can also determine whether the discounting rates differ for rewards of different sizes.

The three experiments differed only in procedural details. In Experiment 1 participants were asked to place their bids without any feedback on the outcomes of any of the auction trials. In Experiment 2 an *indifference-confirmation component* was added in which participants were asked on each trial whether they would rather have their proposed bid or the delayed reward. If participants expressed a preference for either alternative, they were asked to adjust the proposed bid up and down accordingly until they expressed indifference between the bid and the delayed reward. Experiment 3 retained the indifference-confirmation component, but provided half of the participants in each auction with full feedback about all other bids on each trial. These experiments addressed two primary hypotheses: (a) whether the discounting rate observed for the smaller reward would be greater than the discounting rate observed for the larger reward in violation of the rate–amount independence assumption and (b) whether, within each reward size, the hyperbolic function in Equation 2 would better account for the data than the exponential function in Equation 1 in violation of the rate–delay independence assumption.

Apparatus

Participants were seated at computers located in separate testing rooms, and at no time did they interact during the experiment. All responding was accomplished by clicking a mouse on on-screen buttons on the computer's display. A local network connected the computers to a central "auctioneer" computer that controlled the presentation of trials and collected the bid data. The experimenter monitored participants through one-way glass on the doors to the testing rooms. Participants could summon the experimenter at any time during the auction by clicking on a button on the screen, but this occurred rarely after the practice trials.

Procedure

At the time of scheduling over the telephone, participants were told that the study involved bidding in a real auction and that this required bringing $20 cash to the session to use for bidding. They were also told that they would be allowed to bid as little or as much as they wished. No one declined to participate when informed of these requirements.

Instructions were presented on the computer monitor, and participants proceeded through them at their own pace. The instructions described the sealed, second-bid auction in considerable detail and provided examples that illustrated why the optimal bid was the participant's true value of the delayed reward. In order to avoid suggesting any particular rate of discounting for the type of rewards that were used in the experimental trials, the examples in the instructions all involved bidding on a used car worth $4,500. The best strategy for bidding in the auction was summarized in the instructions as follows: "The best strategy is to bid *exactly what the item is worth to you.* The easiest way to decide how much to bid is to ask yourself what is the *most* you would be willing to pay."

When delays to rewards are beyond the end of the experimental session, as they were in these experiments, participants cannot be rewarded for multiple trials without creating dependencies between trials. For example, on Trial 2 participants might try to take into account in their bidding the reward that they were due to receive from Trial 1, and so on, which would hopelessly confound the data on all subsequent trials. For this reason, one trial was selected at random to determine the participant's actual reward. The instructions explained that one of the auction trials would be selected at random at the end of the experiment, that the winner of the auction on that trial would be asked to pay the amount of the second-highest bid on that trial before leaving, and that the experimenter would deliver the delayed reward in cash to the winner in the number of days specified. Participants were told that if they would be out of town on the day the delayed reward was due, the cash would be delivered to them by overnight mail. In order to ensure a preferred outcome on the randomly selected trial, it was in the interest of the participants to treat every trial as though it were the one that would be selected, and they were encouraged to do so by the instructions. When all participants were finished with the instructions, they were provided four practice rounds, both before and after which the experimenter asked each participant whether he or she had any questions.

During the auction a screen display stated, "On this round the item up for auction is," followed by the reward and its delay, for example, "$10.00 in 15 days." Below this, the display read, "The most I would be willing to pay for this item immediately is," followed by a number corresponding to the person's bid. This bid was initialized on each trial at $0, and participants could adjust the bid up and down in 10-cent increments by clicking the mouse on up-and-down arrows, respectively. Bids could not go below $0 nor above the amount of the delayed reward. After a participant settled on a bid, he or she clicked on a "seal bid" button that sent the bid to the auctioneer. In the intertrial interval, participants were provided the message, "Waiting for other players to bid." After all bids were received by the auctioneer, the next delayed reward was presented on the screen.

Two delayed-reward amounts were used: $10 and $20. Both rewards were presented on 15 trials each, using the odd-numbered delays between 1 and 29 days, for a total of 30 auction trials. The rewards alternated between trials, beginning with the $10 reward, and the order of the delays for the two amounts was contrived so that delay did not vary systematically with trial order. All participants received the same ordering of trials. Iterative nonlinear regression analyses were used to fit the discounting functions in Equations 1 and 2 to the 15 bids for each reward amount within subjects. Each analysis yielded an R^2 and a discounting rate parameter (k in Equations 1 and 2) for each amount for each participant. These R^2s and parameters were then used as the data points in matched-pairs t-tests or repeated-measures analyses of variance (ANOVAs). Prior to analysis the R^2s were normalized using a hyperbolic arc-tangent transformation, and the discounting rate parameters were normalized using the natural log. The values reported later are transformed back to their original scale.

At the end of the experiment, participants were asked to report in writing any strategies they used or other information "that would help us understand why you bid the way that you did." After this was completed, the experimenter initiated the selection of one random trial, and a computer display informed participants whether they had won or lost the bid on the selected trial and the amount that the winner would pay for the delayed reward (the second-highest bid).

Experiment 1

Method

The participants were 24 Williams College undergraduates, enrolled in an introductory psychology course, who were recruited with voluntary sign-up sheets offering course credit. Their ages ranged from 18 to 22 years. Ten were men, and 14 were women. The auctions were conducted in seven groups of 3 or 4 participants each. Of the 7 participants who won the bid on the selected trial, 4 paid an average of $4.98 for $10 rewards delayed by an average of 15 days, and 3 paid an average of $17.83 for $20 rewards delayed by an average of 5 days. All delayed rewards were paid in cash on the day specified. The remaining participants did not receive monetary compensation.

Results and Discussion

One participant always bid the amount of the delayed reward, and therefore, her data could not be used to distinguish between different discounting functions. In her post-session report she stated that, for example, the $20 reward in 29 days would still be worth more than $19.90 to her. This represents a discounting rate parameter (hyperbolic or exponential) of less than 0.0002. The following results are for the remaining 23 participants.

Rate–amount independence. Discounting rates decreased as the amount of the delayed reward increased. For 22 of the 23 participants the $10 reward had a higher discounting rate than the $20 reward, $p < .0001$ by sign test, whether hyperbolic or exponential rates were used. The mean hyperbolic rate parameter (k in Equation 2) for the $10 reward was 0.088 ($\pm SE$: 0.065 to 0.117), which was significantly higher than the mean rate parameter for the $20 reward, 0.050 ($\pm SE$: 0.036 to 0.071), $t(22) = 6.79, p < .0001$. The mean exponential rate (k in Equation 1) for the $10 reward was 0.056 ($\pm SE$: 0.043 to 0.073), which was significantly greater than the mean rate for the $20 reward, 0.036 ($\pm SE$: 0.026 to 0.049), $t(22) = 7.00, p < .0001$. Thus, using either the hyperbolic or the exponential function, the discounting rates were inversely related to the amounts of the rewards in violation of the rate–amount independence assumption.

Examples of bids and fitted hyperbolic functions for 3 of the participants, those with discounting rate parameters for the $20 reward at or nearest the 1st, 2nd, and 3rd quartiles, are shown in the upper, middle, and lower panels, respectively, of Figure 2. Overall, the plots of the discounting curves for 19 of the participants were remarkably orderly and showed no systematic departures from the fitted functions. However, the bids for 4 participants revealed discounting curves with unusual shapes. One participant always bid $5 for the $10 reward regardless of delay and

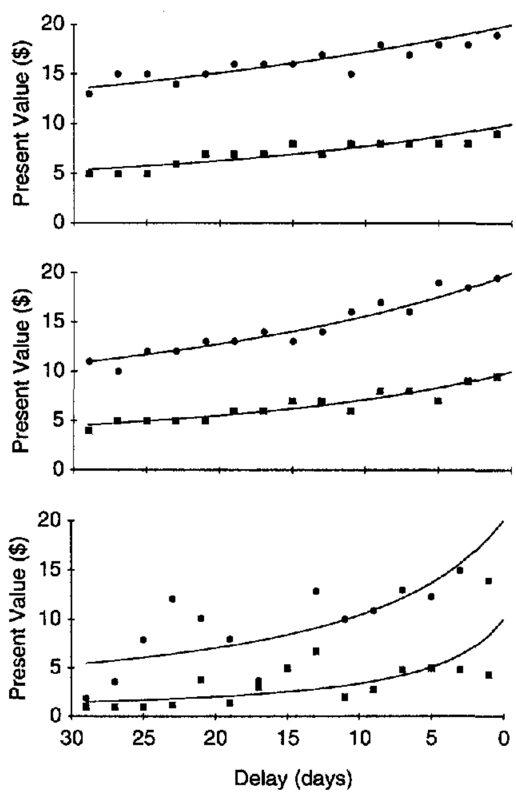

Figure 2. The upper, middle, and lower panels show the data for the participants at or nearest the 1st, 2nd, and 3rd quartiles, respectively, in discounting rate for the $20 reward in Experiment 1. The circles show bids for the $20 reward; the squares show bids for the $10 reward. The curves show the best fitting hyperbolic discounting functions (Equation 2) for each reward size.

always bid \$11 for the \$20 reward regardless of delay. In her self-report she said, "The number of days that the money would be delayed wasn't really an issue for me." A second participant bid about half of the value of the delayed reward up to delays of around 15 days and then dropped off to bids less than \$2 for longer delays, giving his discounting curve a sigmoid shape. Finally, 2 other participants always bid within a narrow band between 10 cents and \$5, indicating an unusually high rate of discounting that took place mostly by a delay of 1 day, with bids declining little thereafter. Neither participant indicated anything in their reports that shed light on their bids. If these 4 participants with unusually shaped discounting curves are excluded, the mean hyperbolic rate parameter for the \$10 reward is 0.060 (±*SE:* 0.048 to 0.076), and the mean rate parameter for the \$20 reward is 0.032 (±*SE:* 0.024 to 0.041). These means remain significantly different, $t(18) = 8.36$, $p < .0001$. Thus, the exclusion of these participants decreases the variance and serves to increase the effect size of the rate–amount relationship, but with or without these participants the qualitative conclusion is the same.

Rate–delay independence. The hyperbolic function fit significantly better than the exponential for both reward amounts. For the \$10 reward, the hyperbolic fit better than the exponential for 19 of the 23 participants, which is significant by sign test, $p = .003$. The mean R^2 for the hyperbolic for the \$10 reward was .961 (±*SE:* .945 to .972) and the mean R^2 for the exponential was .947 (±*SE:* .924 to .964), $t(22) = 4.54$, $p = .0002$.[3] For the \$20 reward, the hyperbolic fit better than the exponential for 20 of the 23 participants, which is significant by sign test, $p = .0005$. The mean R^2 for the hyperbolic for the \$20 reward was .979 (±*SE:* .969 to .986) and the mean R^2 for the exponential was .973 (±*SE:* .958 to .982), $t(22) = 3.24$, $p = .004$.

The data from the 4 participants with unusually shaped discounting curves all favored the hyperbolic function, even though their shapes were clearly not hyperbolic. To ensure that the data from these participants did not account for the relative superiority of the hyperbolic function overall, the analyses were repeated excluding all 4 of these participants. The hyperbolic function still fit significantly better than the exponential for both reward amounts. For the \$10 reward the hyperbolic fit better than the exponential for 15 of the remaining 19 participants, which is significant by sign test, $p = .02$. The mean R^2 for the hyperbolic was .976 (±*SE:* .967 to .982) and the mean R^2 for the exponential was .968 (±*SE:* .955 to .977), $t(18) = 3.62$, $p = .002$. For the \$20 reward the hyperbolic fit better than the exponential for 16 of the remaining 19 participants, which is significant by sign test, $p = .004$. The mean R^2 for the hyperbolic was .988 (±*SE:* .983 to .991) and the mean R^2 for the exponential was .985 (±*SE:* .978 to .990), $t(18) = 2.44$, $p = .025$. Thus, excluding the 4 participants with unusually shaped discounting curves does not alter the conclusion that the hyperbolic function fit better than the exponential function and, therefore, that the discounting rate is inversely related to delay.

Experiment 2

Method

Participants. The participants were 28 people from the Williams College community, including summer students, college staff, and persons unaffiliated with the college, who were recruited through sign-up fliers and a newspaper advertisement. Twelve participants were men, and 16 were women, with ages ranging from 17 to 65 years. Auctions were held in eight groups of 3 or 4 participants each. Of the 8 winners on the selected trials, 5 paid an average of \$6.12 for \$10 rewards delayed by an average of 10 days, and 3 paid an average of \$16.33 for \$20 rewards delayed by an average of 23 days. One participant's delayed reward was delivered by overnight mail to an out-of-town location. All participants were paid at a rate of \$5 per hour for their participation. All delayed rewards were paid in cash.

Procedure. The procedure was identical to that described in the General Methods section except for the addition of an indifference-confirmation component. When the participant had finished adjusting the bid and clicked on the "seal bid" button, a dialog box appeared on the screen that said, "Which would you rather have," followed by a button displaying the amount that the participant bid, another button displaying the delayed reward, and a third button labeled "about equal." For example, if the participant bid \$19 for the \$20 reward in 5 days, the three buttons would read "\$19.00 today," "\$20.00 in 5," and "about equal." If the participant clicked on the amount he or she bid, the computer responded, "You should bid lower." If the participant clicked on the delayed reward, the computer responded, "You should bid higher." The participant was then allowed to readjust the bid. Only when the participant clicked on the "about equal" button was the bid accepted and entered into the auction.

Results and Discussion

Three of the 28 participants in Experiment 2 never bid below the amount of the delayed reward. In their self-reports they all confirmed that their discounting rates were too low for the 29-day delay to matter, corresponding to discounting rates less than 0.0002 per day. The results that follow are for the remaining 25 participants.

Rate–amount independence. As in Experiment 1, the discounting rates decreased as the amount of the delayed reward increased. For 24 of the 25 participants the \$10 reward had a higher discounting rate than the \$20 reward, $p = .0003$ (by sign test), for either the hyperbolic or exponential functions. The mean hyperbolic rate parameter for the \$10 reward was 0.035 (±*SE:* 0.024 to 0.050), which was significantly greater than the mean rate parameter for the \$20 reward, 0.018 (±*SE:* 0.013 to 0.026), $t(24) = 10.50$, $p < .0001$. The mean exponential rate for the \$10 reward was 0.024 (±*SE:* 0.017 to 0.034), which was significantly greater than the mean rate for the \$20 reward, 0.014 (±*SE:*

[3] The raw R^2s reported here are very large because the fitted functions contained no constant term. As is clear from the second panel of Figure 1, exponential and hyperbolic functions are highly correlated and therefore usually account for similar proportions of the variance in discounting data. It is the consistency with which the hyperbolic function provides a relatively better fit that is important here, even though the mean difference in R^2 is small.

0.010 to 0.020), $t(24) = 10.10$, $p < .0001$. Therefore, for either function it is clear that the discounting rate decreased as the amount of the reward increased.

Rate–delay independence. Again, the hyperbolic function fit better than the exponential function for both reward sizes. For the $10 reward the hyperbolic fit better than the exponential for 23 of 25 participants (the remaining 2 were ties to 6 decimal points), $p < .0001$. The mean R^2 for the hyperbolic was .985 ($\pm SE$: .976 to .990) and the mean R^2 for the exponential was .979 ($\pm SE$: .966 to .987), $t(24) = 6.02$, $p < .0001$. For the $20 reward the hyperbolic fit better than the exponential for 22 of the 25 participants (again, with 2 ties), $p = .0002$. The mean R^2 for the hyperbolic was .992 ($\pm SE$: .987 to .995) and the mean R^2 for the exponential was .990 ($\pm SE$: .984 to .994), $t(24) = 3.62$, $p = .001$.

The bids and fitted hyperbolic functions for the 3 participants with discounting rate parameters for the $20 reward at or nearest the 1st, 2nd, and 3rd quartiles, are shown in Figure 3. Only 1 participant gave bids that obviously diverged from any of the fitted functions, discounting the delayed rewards to about one fourth of their value with a 1-day delay and discounting little more for longer delays. This person had by far the lowest R^2s, .500 and .588 for the $10 and $20 rewards, respectively. Even though the hyperbolic function fit better than the exponential for this person, excluding this participant from the analysis made little difference in the results and actually increased the significance of the mean difference between the hyperbolic and exponential fits due to the reduction in error variance.

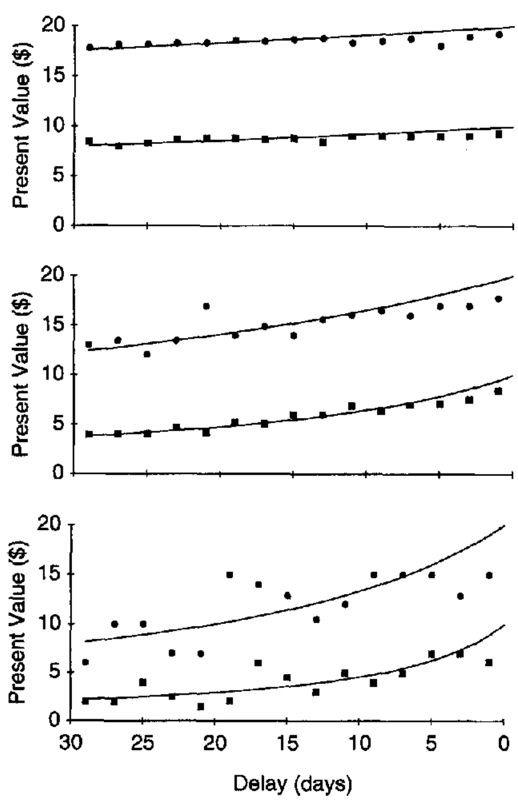

Figure 3. The upper, middle, and lower panels show the data for the participants at or nearest the 1st, 2nd, and 3rd quartiles, respectively, in discounting rate for the $20 reward in Experiment 2. The circles show bids for the $20 reward; the squares show bids for the $10 reward. The curves show the best fitting hyperbolic discounting functions (Equation 2) for each reward size.

The rate of unusually shaped discounting curves in Experiment 2 (3%) is not quite significantly lower than the rate in Experiment 1 (14%), $\chi^2(1) = 2.09$, $p = .07$ (one-tailed). The R^2s in Experiment 2 were also not quite significantly higher than those in Experiment 1: For the $10 reward, $\chi^2(1) = 2.44$, $p = .06$ (one-tailed); for the $20 reward, $\chi^2(1) = 2.10$, $p = .07$ (one-tailed). However, all of these results are consistent in suggesting that the use of indifference confirmation in Experiment 2 may have improved participants' understanding of the properties of the second-bid auction and improved the accuracy of their bids. Given that the two experiments sampled from different populations of participants, this conclusion should be accepted only tentatively and cautiously. However, because there is little cost in adding the indifference-confirmation component, there is no apparent reason not to use it in future research.

Experiment 3

Method

Participants. The participants were 20 people from the Williams College community, including summer students, college staff, and persons unaffiliated with the college, with ages ranging from 17 to 79 years. All were recruited through sign-up fliers and a newspaper advertisement. Of these participants, 7 were men and 13 were women. Auctions were held in five groups of 4 participants each, with 2 participants receiving feedback and 2 participants not receiving feedback in each group. Of the 5 winners on the selected trials, 3 paid an average of $6.80 for $10 rewards delayed by an average of 6 days, and 3 paid an average of $14.90 for $20 rewards delayed by an average of 17 days. All participants were paid at a rate of $5 per hour for their participation. All delayed rewards were paid in cash.

Procedure. The procedure was identical to that used in Experiment 2 except that 2 of the 4 participants in each auction group were given complete feedback about the bidding on each trial. This feedback consisted of the bids of all 4 participants in rank order, with the winner's bid and each participant's own bid (if not the winner) clearly labeled. The participants who received feedback were not aware that other participants were not receiving feedback, and the participants who did not receive feedback were not aware that other participants were receiving feedback.

Results and Discussion

One participant in the feedback condition reported that he believed the experiment was a "scam" to get money from the participants and bid $0 for every trial. The following results are for the remaining 19 participants.

Feedback versus no feedback. Participants in the feedback conditions bid higher, on average, than the participants in the no-feedback condition. Using Equation 2 to generate discounting rate parameters for each participant within reward sizes, the data were analyzed in a two-way ANOVA, with feedback condition as a between-subjects factor and reward size as a within-subjects factor. In the feedback condition the mean discounting rate was 0.016 ($\pm SE$: 0.010 to 0.027), which is significantly lower than the no-feedback condition, 0.096 ($\pm SE$: 0.059 to 0.154), $F(1, 17) = 6.60$, $p = .02$. There was no interaction between feedback condition and reward size, $F(1, 17) = 1.43$, $p = .25$. From participants' self-reports, it appears that a number of people in the feedback condition bid upward to ensure that they won the reward, even when this meant bidding more than the delayed reward alone was worth to them. Two participants explicitly reported increasing their bids in order to win. One participant who consistently came out as the second-highest bid had decided to try to bid as high as possible without winning in order to make the winner pay as much as possible! One participant reported trying to bid low in order to get other bidders to decrease their bids cooperatively.

In all, feedback participants won on 77% of the 150 auction trials across the five groups; if the single group in which a no-feedback participant won the majority of the bids is excluded, the percentage of feedback participants who won the bid increases to 89%. For the $10 reward, no-feedback participants bid $4.66 on average, whereas the feedback participants bid $7.06 on average, $t(17) = 2.56$, $p = .02$. For the $20 reward, no-feedback participants bid $10.18 on average, whereas the feedback participants bid $15.74 on average, $t(17) = 2.68$, $p = .02$. Therefore, it appears that at least a substantial minority of feedback participants strategically bid higher than the amount representing their true discounted values of the delayed rewards, and at least one participant bid low in an effort to reduce bids. This suggests that any benefits of using feedback to help participants understand the second-bid auction procedure are more than outweighed by the resulting inaccuracies in their bids as measures of discounted value alone. Essentially, the feedback introduces competitive and cooperative aspects to outcomes of the auction that change the nature of the item for which the participant is bidding to, for example, a "delayed-money-plus-satisfaction-of-winning-now" reward bundle.

Rate–amount independence. Nonetheless, as in Experiments 1 and 2, the discounting rates decreased as the amount of the delayed reward increased. For 16 of the 19 participants the $10 reward had a higher discounting rate than the $20 reward, $p = .004$ (by sign test), for either the hyperbolic or exponential functions. The mean hyperbolic rate parameter for the $10 reward was 0.052 ($\pm SE$: 0.036 to 0.075), whereas the mean rate parameter for the $20 reward was 0.033 ($\pm SE$: 0.022 to 0.050), $t(18) = 4.25$, $p = .0005$. The mean exponential rate for the $10 reward was 0.035 ($\pm SE$: 0.025 to 0.049), which was significantly greater than the mean rate for the $20 reward, 0.024 ($\pm SE$: 0.016 to 0.034), $t(18) = 4.23$, $p = .0005$. Thus, using either function the discounting rate decreased as the amount of the reward increased.

Rate–delay independence. The results for the comparisons between the hyperbolic and exponential discounting functions replicated those in Experiments 1 and 2. For the $10 reward the hyperbolic function fit better than the exponential function for 17 of 19 participants, $p = .0007$. The mean R^2 for the hyperbolic was .967 ($\pm SE$: .947 to .979) and the mean R^2 for the exponential was .957 ($\pm SE$: .928 to .974), $t(18) = 3.47$, $p = .003$. For the $20 reward the hyperbolic again fit better than the exponential for 17 of the 19 participants, $p = .0007$. The mean R^2 for the hyperbolic was .982 ($\pm SE$: .968 to .990) and the mean R^2 for the exponential was .976 ($\pm SE$: .957 to .987), $t(18) = 3.10$, $p = .006$.

The bids and fitted hyperbolic functions for the 3 participants with discounting rates for the $20 reward at or nearest the 1st, 2nd, and 3rd quartiles, are shown in Figure 4. The plots of the discounting curves showed no systematic irregularities, except, as in Experiment 2, for 1 participant who discounted the delayed rewards to under approximately one fourth of their value with a 1-day delay and discounted little more for longer delays. Again, this participant had by far the lowest R^2s, .183 and .542 for the $10 and $20 rewards, respectively. Excluding this participant from the analysis made no substantive difference in the results.

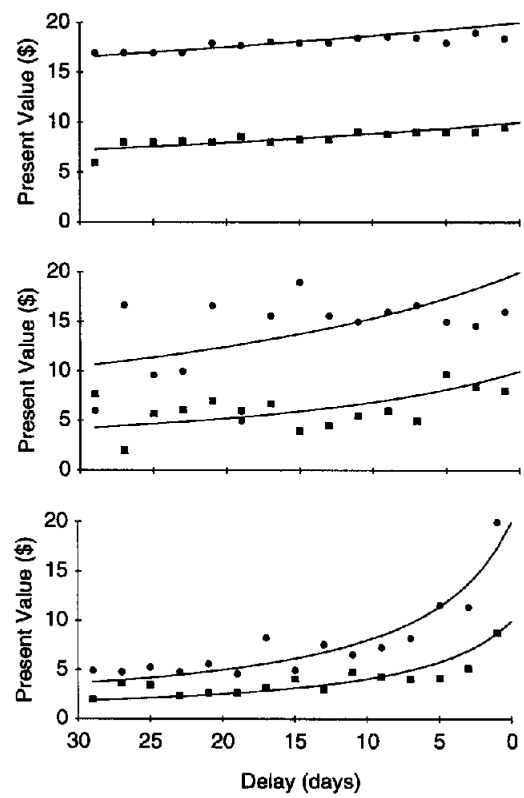

Figure 4. The upper, middle, and lower panels show the data for the participants at or nearest the 1st, 2nd, and 3rd quartiles, respectively, in discounting rate for the $20 reward in Experiment 3. The circles show bids for the $20 reward; the squares show bids for the $10 reward. The curves show the best fitting hyperbolic discounting functions (Equation 2) for each reward size.

Aggregate Data Across Experiments 1, 2, and 3

Figure 5 shows the median bids at each delay for each reward amount across Experiments 1, 2, and 3, excluding the participants who received feedback in Experiment 3. The interquartile ranges are located between the upper and lower bars, and the best fitting hyperbolic function for the median bids is shown by the solid curve. For both the \$10 and \$20 rewards, the hyperbolic function fit the data better than the exponential function, and the hyperbolic discounting rate parameter for the \$20 reward (0.026) was smaller than that for the \$10 reward (0.052). Interestingly, for short delays the fitted hyperbolic curves consistently overestimated the median bids. One possible explanation is that people expect a one-time "premium" for accepting any delay, but then discount smoothly and normatively with additional delays: the "one-period-realization-of-risk" hypothesis (Benzion et al., 1989). Three considerations argue against this explanation. First, this hypothesis failed to account for Benzion et al.'s data. Second, such an explanation would be difficult to reconcile with the observed preference reversals in Kirby and Herrnstein (1995) because it predicts that discounting curves cross as the delay to the smaller reward is increased from zero to the shortest delay period and remain in a fixed relationship thereafter. Across Kirby and Herrnstein's three experiments only 9% of the observed preference reversals occurred by the shortest nonzero delay to the smaller reward. Third, one can recompute the fitted curves by pretending that the delayed reward was equal to the median bid for the 1-day delay, essentially ignoring the one-period decrease that takes place with the initial 1-day delay. Although the difference is smaller in this case, the hyperbolic function still fits better for both rewards, and the rates remain inversely related to amount. Alternatively, one can refit the curves while treating the amount of the delayed reward as a second free parameter. This allows the regression to adjust for any one-period drop. When this is done with the present data, the hyperbolic still fits better for the \$10 reward, and for the \$20 reward the exponential and hyperbolic fits are indistinguishable to four decimal places. For both rewards the discounting rates remain inversely related to amount. In any case, one should be very cautious in interpreting trends found in aggregate data. The same sort of artifactual phenomenon described at the beginning of this article could tend to level these curves. The medians do not correspond to any real bidder, and the overestimates at short delays were not consistently observed in the data of individual participants. Further empirical work is necessary to determine whether the overestimates are consistently found, and if so, whether a function with even more "bend" than Equation 2 is required to account for the data.[4]

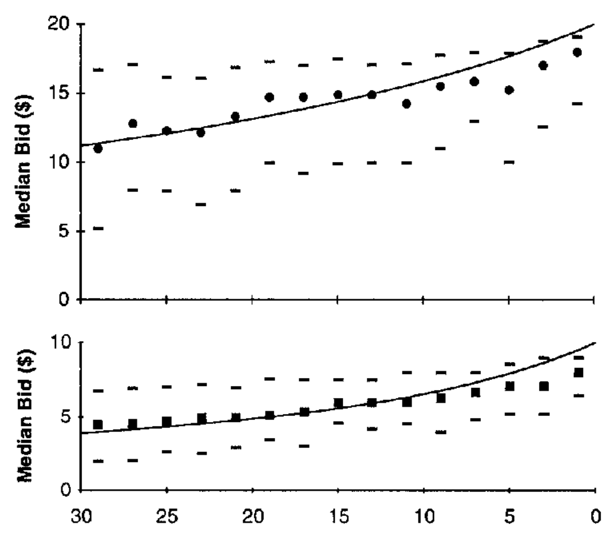

Figure 5. The upper panel (\$20 reward) and lower panel (\$10 reward) show the median bids as a function of delay for all participants in Experiments 1, 2, and 3. The bounds on the interquartile ranges at each delay are indicated by the short horizontal bars. The curves show the best fitting hyperbolic functions (Equation 2) for the median bids.

The top two panels in Figure 6 show the median second-highest bids across the 15 auctions in Experiments 1 and 2, along with their interquartile ranges and best fitting hyperbolic curves. The bottom panel in Figure 6 shows the second-highest bids that would have been observed had all participants in Experiments 1, 2, and 3 (excluding the participants who received feedback) participated in the same auction. That is, these second bids were generated artificially by ignoring groups, which gives an approximation of what the second-highest bids might have been for an auction with 62 participants. The second-highest bids represent the selling prices for the delayed rewards, and it is possible that these could follow the normative model even when the bids themselves do not. However, for the aggregate data in each of the top two panels, the hyperbolic function fit better than did the exponential. In the bottom panel the curves are so flat that the two functions were indistinguishable. Nonetheless, in all three panels the estimated discounting rates were inversely related to reward size.

General Discussion

The results of all three experiments strongly supported both hypotheses: (a) that the delay-discounting rate is inversely related to the amount of the reward (true for 62 of 67 participants) and (b) that the delay-discounting rate is inversely related to the length of the delay to a reward (true for 59 of 67 participants). These results strongly suggest that the rate–delay and rate–amount independence assumptions underlying the normative discounting model do not describe actual discounting. These experiments overcome the limitations of previous research in several ways. First, although the second-bid auction procedure cannot guarantee that participants give their true present values of the delayed re-

[4] See Loewenstein and Prelec (1992) for a generalized hyperbolic function that can bend even more sharply than the one in Equation 2.

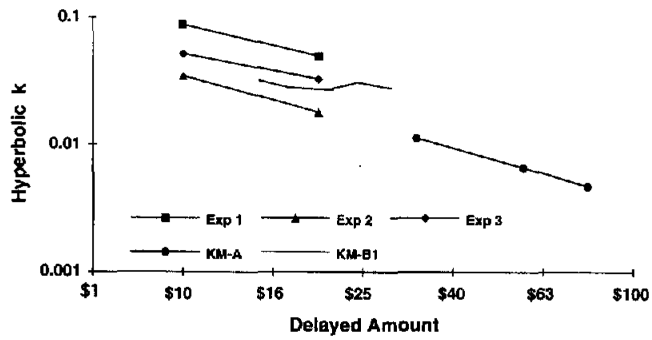

Figure 7. Hyperbolic discounting rate parameters (k in Equation 2) by delayed reward amount for experiments using real rewards. Axes are in $\log_{10}$ coordinates. KM-A = Kirby and Maraković (1996); KM-B1 = Kirby and Maraković (1995), Experiment 1; Exp = Experiment.

wards, it does give them an incentive for doing so as accurately as they can. Second, both violations of the normative assumptions were demonstrated within subjects, which avoids artifacts due to averaging value across subjects. Third, the inverse rate–delay relationship was demonstrated within reward amounts, thereby avoiding any tacit assumption of rate–amount independence. Finally, every participant had real money at stake in their valuations of the delayed rewards.

To explore the consistency of the rate–amount relationship across studies, all of the previous studies that have found an inverse relationship between rate and amount, and that reported data from which this relationship can be estimated, are summarized in Figures 7 and 8. For each study, hyperbolic (Equation 2) discounting rate parameters are shown as a function of amount in log coordinates. The experiments shown in Figures 7 and 8 employed a wide variety of procedures and analytical techniques. Figure 7 shows the discounting rates from the experiments in which subjects received real rewards: Experiments 1, 2, and 3 reported in this paper; the low-bid wins auction in Kirby and Maraković (1995, Experiment 1); and the choice questionnaire study in which participants were rewarded by lottery (Kirby & Maraković, 1996). The slopes of the lines in Figure 7 are −0.82, −0.96, and −0.66 for Experiments 1, 2, and 3 (including feedback participants), respectively, −1.0 for Kirby and Maraković (1996) and −0.10 for Kirby and Maraković (1995).

Figure 8 shows the discounting rates for the studies that have examined the inverse rate–amount relationship using hypothetical rewards (Benzion et al., 1989; Green, Fristoe, et al., 1994; Green, Fry, et al., 1994; Raineri & Rachlin, 1993; Thaler, 1981).[5] For each of these studies, hyperbolic discounting rates were estimated separately within each reward size based on data reported in the published articles. The two studies using questionnaires are shown with dotted lines in Figure 8. Thaler (1981) asked participants to specify a delayed amount of money that would be "just as attractive" as a given immediate reward. The figure shows the estimated discounting rates based on the median amounts specified by participants at each delay (slope = −0.39). Benzion et al. (1989) asked participants to either specify a delayed amount of money that would be equal in value to a given immediate reward ("postpone a receipt" condition; slope = −0.17) or to specify an immediate amount that would be equal to a given delayed reward ("expedite a receipt" condition; slope = −0.11). The rates in Figure 8 are estimated from the mean amounts specified in each condition. In Kirby and Maraković's (1995) Experiment 2 participants were asked to specify the present values of delayed hypothetical rewards, but each trial was presented separately by a computer rather than in a questionnaire. Although the rate–amount relationship was not significant in that study, the slope of the line is −0.19, which is well

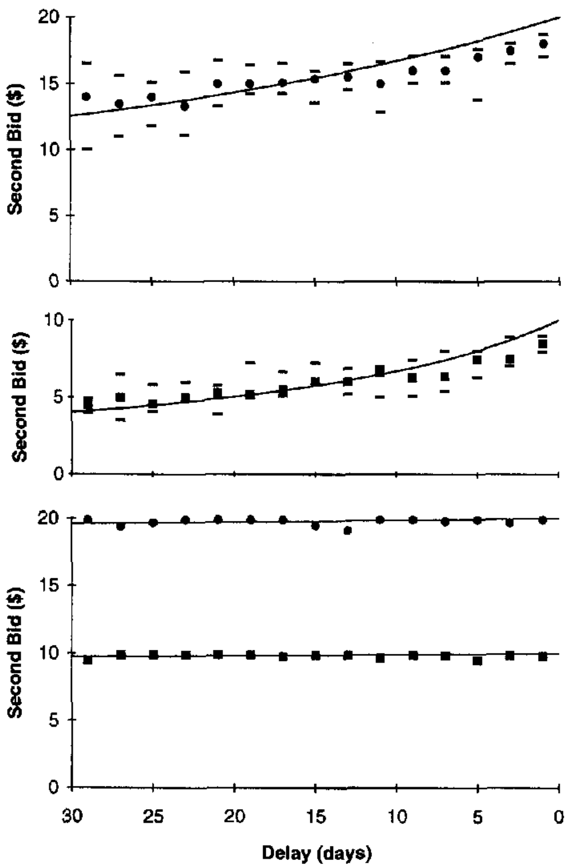

Figure 6. The upper panel ($20 reward) and middle panel ($10 reward) show the median second bids across the 15 auction groups in Experiments 1 and 2 as a function of delay. The bounds on the interquartile ranges are indicated by the short horizontal bars. The curves show the best fitting hyperbolic functions (Equation 2) for the median bids. The bottom panel shows for both rewards the second-highest bids across all participants in all three experiments (excluding participants who received feedback in Experiment 3), as though all 62 bidders participated in the same auction, along with the best fitting hyperbolic functions.

[5] Benzion et al. (1989) and Loewenstein (1987) also found an inverse rate–magnitude relationship for delayed hypothetical losses.

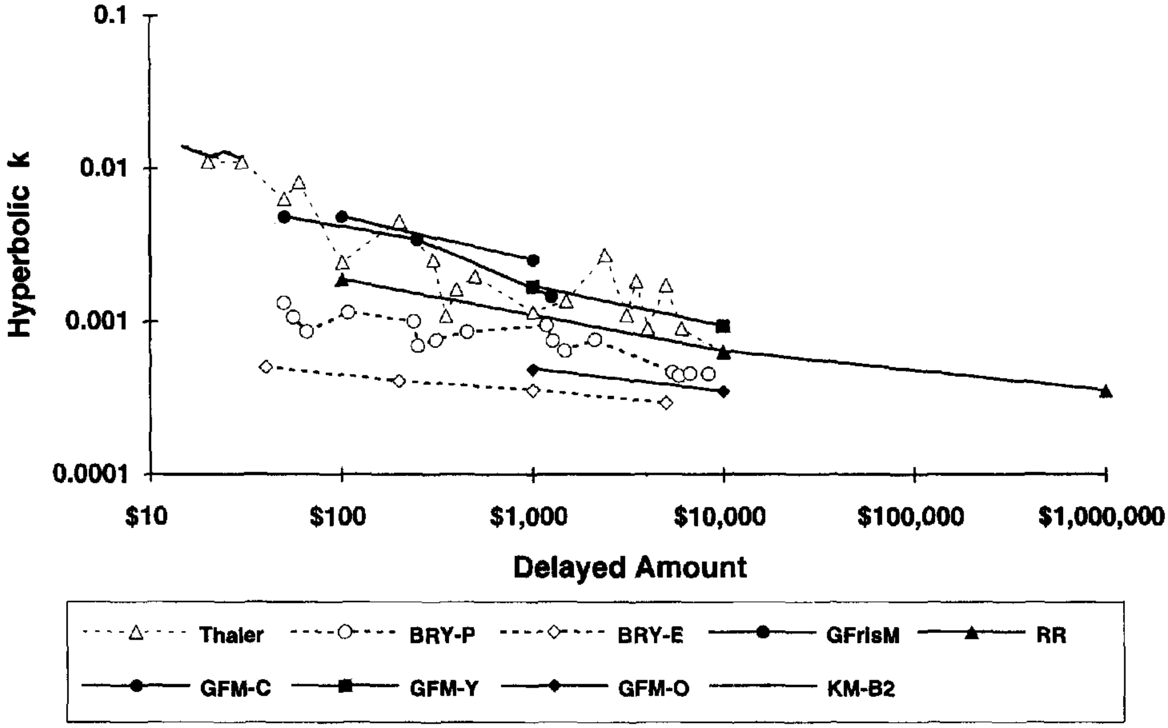

Figure 8. Hyperbolic discounting rate parameters (k in Equation 2) by delayed reward amount for experiments using hypothetical rewards. Axes are in $\log_{10}$ coordinates. Thaler = Thaler (1981); BRY-P = Benzion et al. (1989), "postpone" condition; BRY-E = Benzion et al. (1989), "expedite" condition; GFrisM = Green, Fristoe, et al. (1994); RR = Raineri and Rachlin (1993), Experiment 1; GFM-C, GFM-Y, and GFM-O = Green, Fry, et al. (1994), children, young adults, and older adults, respectively; KM-B2 = Kirby and Maraković (1995), Experiment 2.

within the range of the slopes for the other hypothetical reward studies.

In the remaining three experiments shown in Figure 8 participants were offered hypothetical choices between smaller, earlier rewards and larger, later rewards at different delays. For two of these studies the amounts of the smaller rewards were manipulated according to participants' choices, and indifference points were estimated between the smaller and larger rewards based on participants' preference reversals (Green, Fry, et al., 1994; Raineri & Rachlin, 1993). Raineri and Rachlin (1993) used rewards up to $1 million and delays up to 50 years. However, Raineri and Rachlin's 25- and 50-year data are omitted from the rate estimates in Figure 8 because they reported that those delays "were apparently well beyond subjects' time horizons" (p. 83–84). The figure shows the rates based on the ratios of the immediate to the delayed rewards for the observed preference reversals (slope = −0.23). Green, Fry, et al. (1994), using a similar procedure, tested people in three age groups: children (slope = −0.28), young adults (slope = −0.26), and older adults (slope = −0.14). The points shown in Figure 8 are based on the median immediate amounts at preference reversal for each delayed reward (estimated from Green, Fry, et al., 1994, Figure 1, p. 35). The data are plotted separately for each age group, and consistent with Green, Fry, et al.'s analysis, it can be seen that the discounting rate decreased with age. Finally, Green, Fristoe, et al. (1994) employed a choice procedure that differed from those above in that the delayed and immediate amount pairs were always the same and it was the delays that were adjusted based on participants' choices. The authors entered the median delays at which the participants were indifferent between the smaller and larger rewards into polynomial regressions for each reward pair. For Figure 8 the rate estimates were derived from the fitted constants from those regressions, which represent the delay to larger reward at indifference when the delay to the smaller reward is zero (slope = −0.37).

Two aspects of Figures 7 and 8 are striking: The lines are both remarkably linear and remarkably parallel. It appears that the logs of the hyperbolic discounting rate parameters are approximately linear in log amount. In Figure 7 the one experiment that diverges substantially from this generalization is the Kirby and Maraković (1995) low-bid-wins auction study, which found no significant relationship between discounting rates and amounts, and has a much flatter slope than the other results in Figure 7. As is clear from Figure 7, a floor effect due to overbidding cannot account for the failure to find a rate–amount relationship. The discounting rate parameters observed in that study fall within the typical range. Furthermore, the failure to find a relationship cannot be due to the use of small rewards or a small range of rewards. Both the average size and the range were smaller in Experiments 1, 2, and 3 than in that study. The results of the low-bid-wins auction remain at odds with the results from all other procedures that have been used to examine the rate–amount relationship. In Figure 8, linearity least well describes the rather noisy data from Thaler (1981), but that study had by far the fewest observations per data point in the figure.

The mean slope of the real reward studies is −0.71, or −0.86 excluding Kirby and Maraković (1995). The mean slope of the hypothetical reward studies is −0.24. With the exception of Kirby and Maraković (1995), the slopes for real and hypothetical rewards do not overlap. One possible explanation is that the relationship between log rate parameter and log amount may actually be convex rather than linear. Because the real rewards used were, in most instances, smaller than the hypothetical rewards, they could have larger slopes simply by virtue of being further to the left on the x-axes in Figures 7 and 8. However, the one study that used hypothetical rewards in the range of the real rewards, Kirby and Maraković (1995), had a slope that much more closely resembled the other hypothetical reward studies than the real reward studies. A second possible explanation is that hypothetical rewards may be discounted differently than real rewards. They have lower discounting rate parameters, and those rates decrease less sharply with amount than do real discounting rates. This explanation has some intuitive appeal. Hypothetical rewards may lack the motivational properties of real rewards: People may not be able to fully imagine how much they want a delayed reward until they have an actual opportunity to receive it. This might lead them to say that they would be more willing to wait for a delayed reward than they actually would and that smaller delayed rewards are more valuable than they actually are.

If an approximation that log rate is linear in log amount is accepted, then the hyperbolic discounting rate parameter k in Equation 2 may be replaced by bA^m, where m is the slope in Figures 7 and 8 and b is the y-intercept (note that the y-axis does not cross the x-axis at 0 in Figures 7 and 8). The exponent m can be interpreted as discounting sensitivity to amount, and the multiplier b is now the discounting rate parameter. Substituting into Equation 2 yields

$$V = \frac{A}{1 + bA^mD}.$$

For hypothetical rewards m averages approximately −0.24, and for real rewards m averages approximately −0.8. In the questionnaire experiment using probabilistic rewards, Kirby and Maraković (1996) found that Equation 3 with m equal to −1 fit better across rewards sizes than did Equations 1 or 2. In addition, Green, Fry, et al. (1994) included in Equation 2 an exponent on delay, which they interpret as sensitivity to delay, and found that this sensitivity increased with age. In Figure 8 the slope of the lines for their data decrease with age, suggesting that discounting sensitivity to amount may also decrease with age.

Existing data do not tell us why discounting rates are inversely related to amount. Loewenstein and Prelec (1992) have suggested that the value function for small rewards may be more sharply convex than the value function for large rewards, which would yield higher observed discounting rates for small rewards. For example, the difference in subjective value between \$1 and \$2 may be proportionally less than the difference between \$10 and \$20. If participants were asked to choose between \$1 now and \$2 in a week and also between \$10 now and \$20 in a week, they might choose the \$1 now but choose to wait the week for the \$20. This would make the discounting rate for the \$2 effectively smaller than the rate for \$20. Alternatively, it is possible that it is not the magnitude of the rewards that makes the difference, but rather the absolute magnitude of the difference between the rewards that matters. People might not be willing to wait a week for an extra dollar (whatever the magnitudes of the two rewards), but would be willing to wait for an extra \$10, even when the two cases represent proportionally the same rate of increase. The explanation of this relationship awaits future research.

Two possible limitations on the generality of the results in the experiments reported here deserve mention. First, the probability of receiving a reward affects present value in ways that are qualitatively similar to the effects of delay (King, Logue, & Gleiser, 1992; Mazur, 1993; Rachlin et al., 1991). There are two ways in which the rewards in these experiments were probabilistic:

1. Only a single trial was chosen from all 30 trials to serve as the reward trial. Participants were encouraged to treat each trial as though it were the only trial they faced, and this was the optimal strategy. However, if people were to take the probability that any given trial would be selected into account in their valuations, one could argue that the expected value of the delayed reward on each trial was less than its nonstochastic delay-discounted value. This possibility, however, makes testable predictions that were not borne out in the data. Because participants did not know in advance how many trials they faced, one might predict that the expected values would decrease throughout the session as participants updated their best guesses about the odds of any individual trial becoming the actual reward trial, ultimately declining to about 1/30th of the true discounted value. In fact, bids did not systematically decline across trials, but tended to slightly increase: The average correlations between trial number and bids were .03, .07, and .03 for the three experiments (none of which were significant, all $ts \leq 1$). Furthermore, no participant, except the one in Experiment 3 who bid \$0 every time, gave more than a few bids that were as low as 1/30th of the undiscounted amounts of the delayed rewards. Alternatively, one might argue that participants could have lowered all of their valuations of the delayed reward by some fixed probability adjustment. Fortunately, this would not compromise the main conclusions of the experiments because reducing the expected values by a fixed proportion would neither artifactually favor a hyperbolic function over an exponential nor would it artifactually yield an inverse relationship between rate and amount.

2. The rewards were also probabilistic in the sense that

any delayed reward has a probability less than 1 of actually being received. The experimenter may be dishonest or irresponsible, and either the experimenter or the participant might die, for example, in the delay interval. In principle, the greater the delay, the lower the probability of receiving the reward. However, in the context of these experiments there was little reason for participants to doubt that the rewards would actually be delivered as promised, and no participant expressed this concern during the experiment or in the debriefing statements. Furthermore, even if a participant was concerned about the possibility of not getting a delayed reward, there would be little change in the probability of receiving the reward over the range of delays used. Any plausible change would be much too small to have a substantial impact on choices, given the typical size of the effect of decreasing probability on choice observed in previous experiments with humans (Rachlin et al., 1991).

Second, it should be noted that the conclusions of these experiments are limited to choices between isolated outcomes (Loewenstein & Prelec, 1993; Stevenson, 1993). For example, when people are asked to choose between sequences of hypothetical outcomes, Loewenstein and Prelec (1993) have shown that (a) people prefer sequences that improve over time and (b) people prefer to spread the rewards uniformly over time. Both of these results are in violation of the dictum "the sooner, the better." Loewenstein and Prelec offer a model for preferences over sequences of outcomes in which discounting weight is a function of both the preferred change in the magnitudes of the outcomes over the course of a sequence and the delay discounting of the component outcomes. Equation 3 could be readily incorporated into Loewenstein and Prelec's model as the single-outcome, delay-discounting function.

In conclusion, the results from these experiments clearly suggest that both the rate–amount independence and the rate–delay independence assumptions required by normative discounting models are typically violated by human participants when real money is at stake. The use of money arguably makes for a conservative test of the normative assumptions in that people's experiences with monetary investments and loans, in which interest is compounded exponentially, might enable them to better approximate normative discounting with money than with other types of rewards. Nonetheless, one challenge for future research is to attempt to replicate these results using nonmonetary rewards, such as food (cf. Logue & King, 1991) and durable goods (cf. Kirby & Herrnstein, 1995). Future research also faces the tasks of further specifying the form of discounting functions and discovering effective ways of bringing people's choices more in line with normative prescriptions using self-control strategies (see e.g., Ainslie, 1992; Ainslie & Haslam, 1992; Mischel & Rodriguez, 1993; Mischel, Shoda, & Rodriguez, 1989). After all, behavior consistent with normative discounting should be our goal, even if it is not our norm. The sealed, second-bid auction represents a feasible procedure for assessing present values of delayed rewards and should facilitate further research on delay discounting and impulsiveness.

References

Ainslie, G. W. (1974). Impulse control in pigeons. *Journal of the Experimental Analysis of Behavior, 21*(3), 485–489.

Ainslie, G. (1975). Specious reward: A behavioral theory of impulsiveness and impulse control. *Psychological Bulletin, 82,* 463–496.

Ainslie, G. (1992). *Picoeconomics: The strategic interaction of successive motivational states within the person.* Cambridge, England: Cambridge University Press.

Ainslie, G., & Haendel, V. (1983). The motives of the will. In E. Gottheil, K. Druley, T. Skodola, & H. Waxman (Eds.), *Etiology aspects of alcohol and drug abuse* (pp. 119–140). Springfield, IL: Charles C. Thomas.

Ainslie, G., & Haslam, N. (1992). Self-control. In G. Loewenstein & J. Elster (Eds.), *Choice over time* (pp. 177–209). New York: Russell Sage Foundation.

Ainslie, G., & Herrnstein, R. J. (1981). Preference reversal and delayed reinforcement. *Animal Learning and Behavior, 9,* 476–482.

Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior, 22,* 231–242.

Baum, W. M., & Rachlin, H. (1969). Choice as time allocation. *Journal of the Experimental Analysis of Behavior, 12,* 861–874.

Belke, T. W., Pierce, W. D., & Powell, R. A. (1989). Determinants of choice for pigeons and humans on concurrent-chains schedules of reinforcement. *Journal of the Experimental Analysis of Behavior, 52*(2), 97–109.

Benzion, U., Rapoport, A., & Yagil, J. (1989). Discount rates inferred from decisions: An experimental study. *Management Science, 35*(3), 270–284.

Boehme, R., Blakely, E., & Poling, A. (1986). Runway length as a determinant of self-control in rats. *Psychological Record, 36*(2), 285–288.

Burns, D. J., & Powers, R. B. (1975). Choice and self-control in children: A test of Rachlin's model. *Bulletin of the Psychonomic Society, 5*(2), 156–158.

Chung, S. H., & Herrnstein, R. J. (1967). Choice and delay of reinforcement. *Journal of the Experimental Analysis of Behavior, 10,* 67–74.

Davison, M., & McCarthy, D. (1988). *The Matching Law: A research review.* Hillsdale, NJ: Erlbaum.

Deluty, M. Z. (1978). Self-control and impulsiveness involving aversive events. *Journal of Experimental Psychology: Animal Behavior Processes, 4,* 250–266.

Deluty, M. Z., Whitehouse, W. G., Mellitz, M., & Hineline, P. N. (1983). Self-control and commitment involving aversive events. *Behaviour Analysis Letters, 3,* 213–219.

De Villiers, P. A., & Herrnstein, R. J. (1976). Toward a law of response strength. *Psychological Bulletin, 83,* 1131–1153.

Estes, W. K. (1956). The problem of inference from curves based on group data. *Psychological Bulletin, 53,* 134–140.

Fishburn, P. C., & Rubinstein, A. (1982). Time preference. *International Economic Review, 23,* 677–694.

Flora, S. R. (1995). Molar and molecular contingencies and effects of punishment in a human self-control paradigm. *Psychological Record, 45,* 261–281.

Flora, S. R., & Pavlik, W. B. (1992). Human self-control and the density of reinforcement. *Journal of the Experimental Analysis of Behavior, 57,* 201–208.

Forzano, L. B., & Logue, A. W. (1992). Predictors of adult humans' self-control and impulsiveness for food reinforcers. *Appetite, 19,* 33–47.

Forzano, L. B., & Logue, A. W. (1994). Self-control in adult humans: Comparison of qualitatively different reinforcers. *Learning and Motivation, 25,* 65–82.

Forzano, L. B., & Logue, A. W. (1995). Self-control and impulsiveness in children and adults: Effects of food preferences. *Journal of the Experimental Analysis of Behavior, 64,* 33–46.

Green, L., Fisher, E. B., Perlow, S., & Sherman, L. (1981). Preference reversal and self-control: Choice as a function of reward amount and delay. *Behaviour Analysis Letters, 1,* 43–51.

Green, L., Fristoe, N., & Myerson, J. (1994). Temporal discounting and preference reversals in choice between delayed outcomes. *Psychonomic Bulletin and Review, 1,* 383–389.

Green, L., Fry, A., & Myerson, J. (1994). Discounting of delayed rewards: A life span comparison. *Psychological Science, 5,* 33–36.

Green, L., & Myerson, J. (1993). Alternative frameworks for the analysis of self-control. *Behavior and Philosophy, 21,* 37–47.

Green, L., & Snyderman, M. (1980). Choice between rewards differing in amount and delay: Toward a choice model of self-control. *Journal of the Experimental Analysis of Behavior, 34,* 135–147.

Herrnstein, R. J. (1981). Self-control as response strength. In C. M. Bradshaw, E. Szabadi, & C. F. Lowe (Eds.), *Quantification of steady-state operant behavior* (pp. 3–20). Amsterdam: Elsevier/North-Holland Biomedical Press.

Herrnstein, R. J. (1990). Rational choice theory: Necessary but not sufficient. *American Psychologist, 45,* 356–367.

Holcomb, J. H., & Nelson, P. S. (1992). Another experimental look at individual time preference. *Rationality and Society, 4,* 199–220.

Horowitz, J. K. (1991). Discounting money payoffs: An experimental analysis. In S. Kaish & B. Gilad (Eds.), *Handbook of Behavioral Economics,* (Vol. 2B, pp. 309–324). Greenwich, CT: JAI Press.

Houston, A., & McNamara, J. (1985). The choice of two prey types that minimises the probability of starvation. *Behavioral Ecology and Sociobiology, 17,* 135–141.

Hyten, C., Madden, G. J., & Field, D. P. (1994). Exchange delays and impulsive choice in adult humans. *Journal of the Experimental Analysis of Behavior, 62,* 225–233.

Ito, M. (1985). Choice and amount of reinforcement in rats. *Learning and Motivation, 16,* 95–108.

Ito, M., & Asaki, K. (1982). Choice behavior of rats in a concurrent-chains schedule: Amount and delay of reinforcement. *Journal of the Experimental Analysis of Behavior, 37,* 383–392.

Killeen, P. (1970). Preference for fixed-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior, 14,* 127–131.

Killeen, P. R. (1985). Incentive theory: IV. Magnitude of reward. *Journal of the Experimental Analysis of Behavior, 43,* 407–417.

King, G. R., & Logue, A. W. (1992). Choice in a self-control paradigm: Effects of uncertainty. *Behavioural Processes, 26,* 143–153.

King, G. R., Logue, A. W., & Gleiser, D. (1992). Probability and delay of reinforcement: An examination of Mazur's equivalence rule. *Behavioural Processes, 27,* 125–137.

Kirby, K. N., & Herrnstein, R. J. (1995). Preference reversals due to myopic discounting of delayed reward. *Psychological Science, 6,* 83–89.

Kirby, K. N., & Maraković, N. N. (1995). Modeling myopic decisions: Evidence for hyperbolic delay-discounting within subjects and amounts. *Organizational Behavior and Human Decision Processes, 64,* 22–30.

Kirby, K. N., & Maraković, N. N. (1996). Delay-discounting probabilistic rewards: Rates decrease as amounts increase. *Psychonomic Bulletin & Review, 3,* 100–104.

Koopmans, T. C. (1960). Stationary ordinal utility and impatience. *Econometrica, 28,* 287–309.

Lancaster, K. (1963). An axiomatic theory of consumer time preference. *International Economic Review, 4,* 221–231.

Loewenstein, G. (1987). Anticipation and the valuation of delayed consumption. *Economic Journal, 97,* 666–684.

Loewenstein, G. (1992). The fall and rise of psychological explanations in the economics of intertemporal choice. In G. Loewenstein & J. Elster (Eds.), *Choice over time* (pp. 3–34). New York: Russell Sage Foundation.

Loewenstein, G., & Prelec, D. (1992). Anomalies in intertemporal choice: Evidence and an interpretation. *Quarterly Journal of Economics, 107,* 573–597.

Loewenstein, G. F., & Prelec, D. (1993). Preferences for sequences of outcomes. *Psychological Review, 100,* 91–108.

Logan, F. A. (1960). *Incentive.* New Haven, CT: Yale University Press.

Logan, F. A. (1965). Decision making by rats: Delay versus amount of reward. *Journal of Comparative and Physiological Psychology, 59,* 1–12.

Logan, F. A., & Spanier, D. (1970). Chaining and nonchaining delay of reinforcement. *Journal of Comparative and Physiological Psychology, 72,* 98–101.

Logue, A. W. (1988). Research on self-control: An integrating framework. *Behavioral and Brain Sciences, 11,* 665–679.

Logue, A. W., Forzano, L. B., & Tobin, H. (1992). Independence of reinforcer amount and delay: The generalized matching law and self-control in humans. *Learning and Motivation, 23,* 326–342.

Logue, A. W., & King, G. R. (1991). Self-control and impulsiveness in adult humans when food is the reinforcer. *Appetite, 17,* 105–120.

Logue, A. W., King, G. R., Chavarro, A., & Volpe, J. S. (1990). Matching and maximizing in a self-control paradigm using human subjects. *Learning and Motivation, 21,* 340–368.

Logue, A. W., Peña-Correal, T. E., Rodriguez, M. L., & Kabela, E. (1986). Self-control in adult humans: Variation in positive reinforcer amount and delay. *Journal of the Experimental Analysis of Behavior, 46,* 159–173.

Logue, A. W., Rodriguez, M. L., Peña-Correal, T. E., & Mauro, B. C. (1984). Choice in a self-control paradigm: Quantification of experience-based differences. *Journal of the Experimental Analysis of Behavior, 41,* 53–67.

Machina, M. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature, 27,* 1622–1668.

Mazur, J. E. (1984). Tests of an equivalence rule for fixed and variable reinforcer delays. *Journal of Experimental Psychology: Animal Behavior Processes, 10,* 426–436.

Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. In M. L. Commons, J. E. Mazur, J. A. Nevin, & H. Rachlin (Eds.), *Quantitative analyses of behavior: The effect of delay and of intervening events on reinforcement value* (Vol. 5, pp. 55–73). Hillsdale, NJ: Erlbaum.

Mazur, J. E. (1993). Predicting the strength of a conditioned reinforcer: Effects of delay and uncertainty. *Current Directions in Psychological Science, 2,* 70–74.

Mazur, J. E. (1994). Effects of intertrial reinforcers on self-control choice. *Journal of the Experimental Analysis of Behavior, 61,* 83–96.

Mazur, J. E., & Herrnstein, R. J. (1988). On the functions relating

delay, reinforcer value, and behavior. *Behavioral and Brain Sciences, 11,* 690–691.
Meyer, R. F. (1976). Preferences over time. In R. L. Keeney & H. Raiffa (Eds.), *Decisions with multiple objectives: Preferences and value tradeoffs* (pp. 473–514). New York: Wiley.
Millar, A., & Navarick, D. J. (1984). Self-control and choice in humans: Effects of video game playing as a positive reinforcer. *Learning and Motivation, 15,* 203–218.
Mischel, W., & Rodriguez, M. L. (1993). Psychological distance in self-imposed delay of gratification. In R. R. Cocking & K. A. Renniger (Eds.), *The development and meaning of psychological distance* (pp. 109–121). Hillsdale, NJ: Erlbaum.
Mischel, W., Shoda, Y., & Rodriguez, M. L. (1989). Delay of gratification in children. *Science, 244,* 933–938.
Navarick, D. J. (1982). Negative reinforcement and choice in humans. *Learning and Motivation, 13,* 361–377.
Navarick, D. J., & Fantino, E. (1976). Self-control and general models of choice. *Journal of Experimental Psychology: Animal Behavior Processes, 2,* 75–87.
Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review, 101,* 53–79.
Pollak, R. A. (1968). Consistent planning. *Review of Economic Studies, 35,* 201–208.
Prelec, D. (1989). *Decreasing impatience: Definition and consequences* (Working Paper No. 90-015). Boston Harvard Business School.
Rachlin, H. (1974). Self-control. *Behaviorism, 2,* 94–107.
Rachlin, H., & Green, L. (1972). Commitment, choice and self-control. *Journal of the Experimental Analysis of Behavior, 17,* 15–22.
Rachlin, H., Raineri, A., & Cross, D. (1991). Subjective probability and delay. *Journal of the Experimental Analysis of Behavior, 55,* 233–244.
Raineri, A., & Rachlin, H. (1993). The effect of temporal constraints on the value of money and other commodities. *Journal of Behavioral Decision Making, 6,* 77–94.
Rodriguez, M. L., & Logue, A. W. (1986). Independence of the amount and delay ratios in the generalized matching law. *Animal Learning and Behavior, 14,* 29–37.
Rodriguez, M. L., & Logue, A. W. (1988). Adjusting delay to reinforcement: Comparing choice in pigeons and humans. *Journal of Experimental Psychology: Animal Behavior Processes, 14,* 105–117.
Solnick, J. V., Kannenberg, C. H., Eckerman, D. A., & Waller, M. B. (1980). An experimental analysis of impulsivity and impulse control in humans. *Learning and Motivation, 11,* 61–77.
Stevenson, M. K. (1993). Decision making with long-term consequences: Temporal discounting for single and multiple outcomes in the future. *Journal of Experimental Psychology: General, 122,* 3–22.
Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies, 23,* 165–180.
Thaler, R. (1981). Some empirical evidence on dynamic inconsistency. *Economic Letters, 8,* 201–207.
Tobin, H., Chelonis, J. J., & Logue, A. W. (1993). Choice in self-control paradigms using rats. *Psychological Record, 43,* 441–453.
Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance, 16,* 8–37.
White, K. G., & Pipe, M.-E. (1987). Sensitivity to reinforcer duration in a self-control procedure. *Journal of the Experimental Analysis of Behavior, 48,* 235–249.
Winston, G. C., & Woodbury, R. G. (1991). Myopic discounting: Empirical evidence. In S. Kaish & B. Gilad (Eds.), *Handbook of Behavioral Economics* (Vol. 2B, pp. 325–342). Greenwich, CT: JAI Press.

Received September 8, 1995
Revision received March 12, 1996
Accepted April 25, 1996 ■